

Learning-based Dynamic Pinning of Parallelized Applications in Many-Core Systems ^{*}

Georgios C. Chasparis¹, Vladimir Janjic², Michael Rossbory¹, and Kevin Hammond²

¹ Software Competence Center Hagenberg GmbH, Softwarepark 21, A-4232 Hagenberg, Austria,

`{georgios.chasparis,michael.rossbory}@scch.at`

² School of Computer Science, University of St Andrews, Scotland, UK, `vj32@st-andrews.ac.uk, kevin@kevinhammond.net`

Abstract. This paper introduces a reinforcement-learning based resource allocation framework for dynamic placement of threads of parallel applications to Non-Uniform Memory Access (NUMA) many-core systems. We propose a two-level learning-based decision making process, where at the first level each thread independently decides on which group of cores (NUMA node) it will execute, and on the second level it decides to which particular core from the group it will be pinned. Additionally, a novel performance-based learning dynamics is introduced to handle measurement noise and rapid variations in the performance of the threads. Experiments on a 24-core system show the improvement of up to 16% in the execution time of parallel applications under our framework, compared to the Linux operating system scheduler.

1 Introduction

Resource allocation has become an indispensable part of the design of any engineering system that consumes resources, such as electricity power in home energy management [1], access bandwidth and battery life in wireless communications [8] and computing bandwidth under different settings [2, 3, 5]. When resource allocation is performed online and the number, arrival and departure times of the tasks are not known a priori (as in the case of CPU bandwidth allocation), the role of a resource manager (RM) is to guarantee an *efficient* operation of all tasks by appropriately distributing resources. which requires the formulation of a centralized optimization problem (e.g., mixed-integer linear programming formulations [2]), which further requires information about the specifics of each task (i.e., application details). Such information may not be available to neither

^{*} This work has been supported by the European Union grant EU H2020-ICT-2014-1 project RePhrase (No. 644235). It has also been partially supported by the Austrian Ministry for Transport, Innovation and Technology, the Federal Ministry of Science, Research and Economy, and the Province of Upper Austria in the frame of the COMET center SCCH.

the RM nor the task itself. Given the difficulties involved in the formulation of centralized optimization problems in resource allocation, not to mention their computational complexity, feedback from the running tasks in the form of performance measurements may provide valuable information for the establishment of efficient allocations. Such (feedback-based) techniques have recently been considered in several scientific domains [4, 5].

This paper proposes a distributed learning scheme specifically tailored for addressing the problem of dynamically assigning/pinning threads of a parallelized application to the available processing units. The proposed scheme is flexible enough to incorporate alternative optimization criteria. In particular, we demonstrate its utility in maximizing the average processing speed of the overall application, which under certain conditions also imply shorter completion time. The proposed scheme also reduces computational complexity usually encountered in centralized optimization problems, while it provides an adaptive response to the variability of the provided resources. This paper extends prior work of the authors [6, 12]. Compared to that work, this paper makes the following research contributions:

- (C1) We propose a novel two-level scheduling process that is more appropriate for Non-Uniform Memory Access (NUMA) architectures. At the higher level, the scheduler decides on which NUMA node each thread should be assigned, while at the lower level the scheduler decides on which CPU core (within that NUMA node) to execute the thread.
- (C2) We propose a novel learning dynamics motivated by *aspiration learning* for making decisions at the higher level.
- (C3) We demonstrate the efficiency of the proposed approach on an application borrowed from the evolutionary computing domain.

The paper is organized as follows. Section 2 discusses related work. Section 3 describes the problem formulation and contributions of the paper. Section 4 presents the main features of the proposed Dynamic Scheduler. Section 5 presents experiments of the proposed resource manager in a many-core Linux platform and comparison tests with the operating system’s response. Finally, Section 6 presents concluding remarks and future work.

2 Related work

Prior work has demonstrated the importance of thread-to-core bindings in the overall performance of a parallelized application. For example, [9] describes a tool that checks the performance of each of the available thread-to-core bindings and searches for an optimal placement. Unfortunately, the *exhaustive-search* type of optimization that is implemented may prohibit runtime implementation. Borguedis et al. [4] combine the problem of thread scheduling with scheduling hints related to thread-memory affinity issues. A similar scheduling policy is also implemented by [11].

The resource allocation problem that we are considering can be seen as centralized optimization problem, where n threads need to be mapped to m processing elements in the optimal way. To reduce this potentially very large search space (m^n possible different allocations), distributed or game-theoretic optimizations have been attempted in the past for related problems, including cooperative game formulation for allocating bandwidth in grid computing [13], the non-cooperative game formulation in the problem of medium access protocols in communications [14] or for allocating resources in cloud computing [15]. These approaches significantly reduce the computation complexity and also allow for the development of online selection rules where tasks/agents make decisions often using current observations of their *own* performance.

Contrary to this line of research, our recent work [6, 12] proposed a dynamic (learning-based) scheme for optimally allocating threads of a parallelized application into a set of available CPU cores. It requires a minimum information exchange, that is only the performance measurements collected from each running thread, and it has linear complexity with the number of running threads. Furthermore, it is flexible enough to accommodate alternative optimization criteria depending on the available performance counters (e.g., average processing speed). It was shown both analytically and through experiments under the Linux operating system, that the proposed methodology learns a locally-optimal allocation, which under certain conditions also corresponds to the global optimum. However, one potential drawback was the fact that no special consideration was taken upon the possible *non-uniform memory access* (NUMA) architectures, as it did not distinguish between moving a thread to a “local” (within the same NUMA node) and “remote” (from the different NUMA node) core.

3 Problem Formulation and Objective

Let a parallel application comprise n threads, $\mathcal{I} = \{1, 2, \dots, n\}$, that need to be pinned to a set \mathcal{J} of available CPU cores, $\mathcal{J} = \{1, 2, \dots, m\}$. We denote the *assignment* of a thread i to the set of available CPU cores by $\alpha_i \in \mathcal{A}_i \equiv \mathcal{J}$, i.e., α_i denotes the CPU core to which the thread is assigned to. Let also $\alpha = \{\alpha_i, i \in \mathcal{I}\}$ denote the *assignment profile*.

The Resource Manager (RM) is responsible for the assignment of threads to CPU cores. It periodically checks the performance of a thread and makes decisions about its placement in the next scheduling iteration, so that a (user-specified) objective is maximized. ***For the remainder of the paper***, we will assume that: a) The internal properties and details of the threads are not known to the RM. Instead, the RM may only have access to measurements related to their performance (e.g., their processing speed); b) Threads may not be idled or postponed. Instead, the goal of the RM is to assign the *currently* available resources to the *currently* running threads.

Static optimization and issues Let $v_i = v_i(\alpha, w)$ denote the processing speed of thread i which depends on both the assignment profile α , as well as exogenous

parameters aggregated within w that summarizes, for example, the impact of other applications running on the same platform (*disturbances*). The centralized objective that we consider then takes the following form:

$$\max_{\alpha \in \mathcal{A}} f(\alpha, w). \quad (1)$$

In this paper, the centralized objective corresponds to the average processing speed of the running threads, i.e.,

$$f(\alpha, w) \doteq \sum_{i=1}^n v_i/n. \quad (2)$$

Any solution to the optimization problem (1) corresponds to an *efficient assignment*. However, there are two practical issues when posing an optimization problem in this form: a) the function $v_i(\alpha, w)$ is unknown and it may only be evaluated through measurements of the processing speed, denoted \tilde{v}_i ; and, b) the exogenous disturbances $w = (w_1, \dots, w_m)$ are unknown and may vary with time, thus the optimal assignment may not be fixed with time.

Measurement- or learning-based optimization We wish to address a *static* objective of the form (1) through a *measurement- or learning-based* optimization approach. According to such approach, the RM reacts to measurements of the objective function $f(\alpha, w)$, periodically collected at time instances $k = 1, 2, \dots$ and denoted by $\tilde{f}(k)$. For example, in the case of objective (2), the measured objective takes on the form $\tilde{f}(k) \doteq \sum_{i=1}^n \tilde{v}_i(k)/n$. Given these measurements and the current assignment $\alpha(k)$ of resources, the RM selects the next assignment of resources $\alpha(k+1)$ so that the measured objective approaches the true optimum of the unknown performance function $f(\alpha, w)$. In other words, the RM employs an update rule of the form:

$$\{(\tilde{v}_i(1), \alpha_i(1)), \dots, (\tilde{v}_i(k), \alpha_i(k))\}_i \mapsto \{\alpha_i(k+1)\}_i \quad (3)$$

according to which prior pairs of measurements and assignments for each thread i are mapped into a new assignment $\alpha_i(k+1)$ that will be employed during the next evaluation interval.

Multi-level decision-making and actuation Recent work by the authors [6, 12] has demonstrated the potential of learning-based optimization of the CPU affinities. However, when an application runs on a Non-Uniform Memory Access (NUMA) machine, additional information can be exploited to enhance scheduling of a parallelized application. To this end, a multi-level decision-making and actuation process is considered. We extend the PaRLSched dynamic scheduler presented in [6, 12] by introducing two nested decision processes depicted in Figure 1. At the *higher level*, the performance of a thread is evaluated with respect to its own prior history of performances, and decisions are taken with

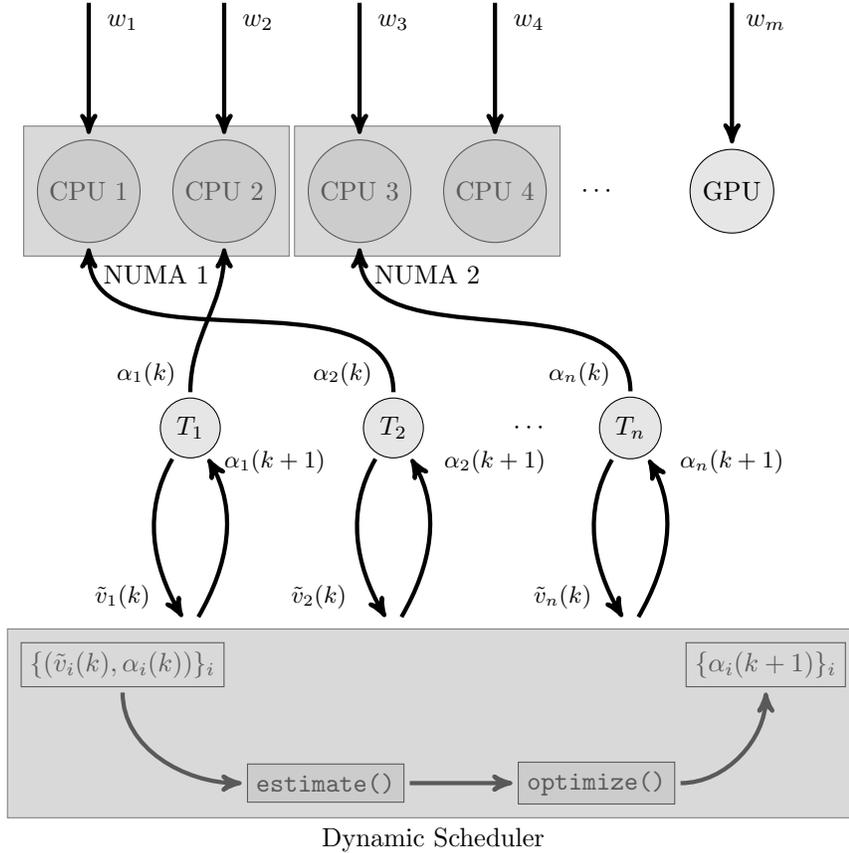


Fig. 1. Schematic of a multi-layer *dynamic* resource allocation framework.

respect to its NUMA placement (possibly involving memory affinities). At the *lower level*, the performance of a thread is evaluated with respect to its own prior history of performances, and decisions are taken with respect to its CPU placement (within the selected NUMA node).

4 Dynamic Scheduler

In this section, we provide a detailed description of the new features of the updated PaRLSched Dynamic Scheduler.

Advancement of architecture According to the new architecture, the user may define the resources to be optimized as well as the corresponding methods used for establishing predictions and for computing optimal allocations. In particular, the initialization of the scheduler accepts the following parameters.

```

RESOURCES={"NUMA_BANDWIDTH", "NUMA_MEMORY"}
OPT_CRITERIA={"PROCESSING_SPEED", "PROCESSING_SPEED"}
RESOURCES_EST_METHODS={"RL", "RL"}
RESOURCES_OPT_METHODS={"RL", "AL"}

```

In the above example, we have defined two distinct resources to be optimized, namely "NUMA_BANDWIDTH", which refers to the placement of threads to specific NUMA nodes, and "NUMA_MEMORY", which refers to the placement/binding of the thread's memory pages into specific NUMA nodes. For each one of the resources to be optimized, there might be alternative optimization criteria, which summarize our objectives for guiding the placements. The selection of the optimization criteria is open-ended and directly depends upon the available performance metrics. In parallel to the selection of the optimized resources, we need to also define the corresponding "methods" for establishing predictions (which are used under the `estimate()` part of the Dynamic Scheduler, Figure 1). Similarly, we may also define the corresponding "methods" for the computation of the next placements (which are used under the `optimize()` part of the Dynamic Scheduler, Figure 1). For example, we may use the Reinforcement-Learning (RL) selection criterion (cf., [6]) for both estimation and optimization, or the Aspiration-Learning (AL) criterion for optimization (described in a forthcoming section).

Hierarchical structure The Dynamic Scheduler has been redesigned so that it accepts a nested description of resources as it is evident in NUMA architectures. The depth of such description is not limited, although in the current implementation we have been experimenting with a single type of a child resource. An example of how child resources can be defined is as follows.

```

CHILD_RESOURCES={"CPU_BANDWIDTH", "NULL"}
CHILD_OPT_CRITERIA={"PROCESSING_SPEED", "PROCESSING_SPEED"}
CHILD_RESOURCES_EST_METHODS={"RL", "RL"}
CHILD_RESOURCES_OPT_METHODS={"AL", "AL"}

```

We have defined a child resource for the NUMA bandwidth resources, which corresponds to a CPU-based description of the placement. For each one of the child resources, we may define separate estimation and optimization methods. Note that the decisions and algorithms over child resources may not coincide with the corresponding methods applied for the case of the original resources.

Aspiration-learning-based dynamics We propose a class of learning dynamics that provide: a) fast response to rapid performance variations, and b) varying exploration rate. Under the reinforcement learning scheme of [6], *exploration* of new affinities is performed with fixed probability and independently of the current performance of the application. It would have been desirable though that the size of exploration varies with the current performance (e.g., larger exploration rate under low performances). To this end, we developed a novel learning scheme that is based upon the notions of benchmark actions/performances and bears similarities with the so-called *aspiration learning*. At each iteration k , the following steps are executed:

1. **Performance update.** At time k , update the (discounted) running average performance of the thread (with respect to the optimized resource), denoted $\bar{v}_i(k)$, as follows:

$$\bar{v}_i(k+1) = \bar{v}_i(k) + \epsilon \cdot [v_i(k) - \bar{v}_i(k)], \quad (4)$$

where $v_i(k)$ is the current measurement of the processing speed of thread i .

2. **Benchmark update.** Define the *upper benchmark performance* $\bar{b}_i(k)$ as follows:

$$\bar{b}_i(k) = \begin{cases} \bar{v}_i(k), & \text{if } \bar{v}_i(k) \geq \bar{b}_i(k-1) \\ \bar{b}_i(k-1), & \text{if } \underline{b}_i(k-1) < \bar{v}_i(k) < \bar{b}_i(k-1) \\ \eta \underline{b}_i, & \text{else,} \end{cases} \quad (5)$$

for some constant $\eta > 1$. The *lower benchmark performance* is defined as follows:

$$\underline{b}_i(k) = \begin{cases} \bar{v}_i(k), & \text{if } \bar{v}_i(k) \leq \underline{b}_i(k-1) \\ \underline{b}_i(k-1), & \text{if } \underline{b}_i(k-1) < \bar{v}_i(k) \leq \bar{b}_i(k-1) \\ (1/\eta) \bar{b}_i, & \text{else.} \end{cases} \quad (6)$$

3. **Action update.** Given the current benchmarks and performance, a thread i selects actions according to the following rule:
 - (a) if $\bar{v}_i(k) < \underline{b}_i(k)$, i.e., if the current average performance is smaller than the lower benchmark performance, then thread i will perform a random switch to an alternative selection according to a uniform distribution.
 - (b) if $\underline{b}_i(k) \leq \bar{v}_i(k) < \bar{b}_i(k)$, then each thread i will keep playing the same action with high probability and experiment with any other action with a small probability $\lambda > 0$.
 - (c) if $\bar{v}_i(k) \geq \bar{b}_i(k)$, i.e., if the current average performance is larger than the upper benchmark performance, then thread i will keep playing the same action.

It is important to note that the above learning scheme will react immediately to a rapid drop in the performance (thus, we indirectly increase the response time to large performance variations). Furthermore, the exploration rate is considerably large under such rapid drops in the performance, and quite small ($\lambda > 0$) when performance remains within the safe region (i.e., between $\underline{b}_i(k)$ and $\bar{b}_i(k)$).

5 Experiments

In this section, we present an experimental study of the proposed framework. Experiments were conducted on `28xIntel@Xeon@CPU E5-2650 v3 2.30 GHz` running Linux Kernel 64bit 3.13.0-43-generic. The cores are divided into two NUMA nodes (Node 1: 0-13 CPU cores, Node 2: 14-27 CPU cores). The experiments were conducted in scenarios where the availability of resources (CPU

cores) may vary over time. We compared the overall performance (in terms of processing speed of threads and completion time of an application) of the framework with that of the OS scheduler. The application that we used was Ant Colony Optimisation (ACO) [7], a metaheuristics used for solving NP-hard combinatorial optimization problems, applied here to the Single Machine Total Weighted Tardiness Problem (SMTWTP). This gives us both computationally- and data-intensive application, where potentially large portions of memory are used for storing and accessing data (making considerations of where the application data is, in terms of NUMA nodes, non-trivial) and where each thread performs large amount of computation over this data. More details about the application can be found in [12].

Experimental setup. For the purpose of reducing the number of experiments, we have fixed the scheduling interval to be 0.2 sec, which is also the interval in which the RM collects measurements of the total instructions per sec (using the PAPI library [10]) for one of the threads separately. This is used as an estimate of the processing speed of each thread. Pinning of threads to CPU cores is achieved through the `sched.h` library (in particular, the `pthread_setaffinity_np` function). In all of the experiments, the RM is executed by the master thread of an application, which is always running in a fixed CPU core (usually the first available CPU core of the first NUMA node). In Table 1, we provide an overview of the conducted experiments with the ACO case study. To evaluate the dynamic scheduler under different quantity of exogenous interferences, we consider four main sets of experiments (A, B, C, and D), where each set differs in the amount of provided resources and their temporal availability. In particular, under the *non-uniform CPU availability* condition, other applications occupy a constant number of the available CPU cores throughout the whole duration of the experiment. On the other hand, under the *time-varying CPU availability* condition, other applications occupy a non-constant part of the available bandwidth (i.e., exogenous applications start running 1 min after the beginning of the experiment). In both the *non-uniform CPU availability* and the *time-varying CPU availability* case, the exogenous disturbances (other applications) comprise computational tasks often equally distributed among the available CPU cores. In the experiment sets A, B, and C, these exogenous applications occupy the first 6 CPU cores of both NUMA nodes. Furthermore, in the set D, we alternate the exogenous interferences between the two NUMA nodes. Our goal is to investigate the effect of the (stack) memory of threads in the overall performance of the application.

5.1 Thread Pinning

Table 2 shows the execution times of the ACO application under OS and PaRLSched scheduler under small, medium and large availability of resources. We can observe the following:

- *Completion time.* For experiments A and B (small and medium amount of resources available), we can observe that under small or no interference

Exp.	Threads	Cores	Amount resources/CPU availability
A.1	40	(8,2)	Small/Uniform
A.2	40	(8,2)	Small/Non-uniform
A.3	40	(8,2)	Small/Time-varying
B.1	40	(8,8)	Medium/Uniform
B.2	40	(8,8)	Medium/Non-uniform
B.3	40	(8,8)	Medium/Time-varying
C.1	40	(12,12)	Large/Uniform
C.2	40	(12,12)	Large/Non-uniform
C.3	40	(12,12)	Large/Time-varying
D	40	(6,6)	Time-varying CPU availability alternating between NUMA nodes.

Table 1. Brief description of the ACO experiments. The tuple in the third column shows the number of CPU cores used in each of the two NUMA nodes of the experimental platform.

Exp/ Time(s)	OS			PaRLSched			Diff. (%)
	Mean	Dev	Avg. Spd.	Mean	Dev	Avg. Spd.	
A.1	1065.05	7.68	13.37	1075.48	6.45	14.87	-0.9
A.2	1752.46	14.00	8.54	1455.92	22.8	9.82	+16.92
A.3	1459.18	9.42	10.29	1402.00	4.06	10.41	+3.91
B.1	673.09	5.69	21.26	699.16	10.24	22.16	-3.87
B.2	1106.36	16.71	12.73	1041.33	16.71	14.94	+5.87
B.3	1066.18	0.88	13.37	1019.11	8.39	14.87	+4.41
C.1	455.87	5.08	31.90	496.26	5.08	33.46	-8.85
C.2	659.78	27.45	21.57	688.80	18.66	24.15	-4.39
C.3	659.35	3.62	21.82	676.03	7.72	23.72	-2.52

Table 2. Completion times of OS and PaRLSched scheduling for A-C experiments ($\epsilon = 0.3/\bar{v}_i/10^8$, $\lambda = 0.1/\bar{v}_i/10^8$). For PaRLSched and OS schedulers, we show the mean execution time of the application, mean deviation (in seconds) and average processing speed of thread (in 10^8 instructions per second).

(experiments A.1 and B.1), the operating system outperforms the PaRLSched scheduler slightly. This is due to the fact that Linux scheduler is utilising internal load balancing of threads between cores, which has notable effect on the execution time when there is not much external interference (in terms of additional running applications). For the PaRLSched scheduler, this benefit is lost due to the explicit pinning of threads to cores, which prevents the operating system to balance the load between cores. This is also a reason why in the set C of experiments the OS scheduler consistently outperforms the PaRLSched scheduler, as there is a large number of cores available and the actual interference by the external application is much lower than in A and B. As a note, in the experiment C only 50% of cores are occupied by the

external application, as opposed to 80% and 75% in the experiments A and B. However, when the interference by the external application is higher (A.2, A.3, B.2 and B.3), we can see that the PaRLSched scheduler outperforms the OS up to 16.92%. This is because under the PaRLSched scheduling, each thread individually adapts to the changes in load. Therefore, for this use case, we can conclude that the higher the load of the system, the bigger the improvement observed under the PaRLSched scheduler.

- *Optimization criterion.* When it comes to measuring how successful the PaRLSched scheduler is in optimizing the scheduling criterion (maximum average processing speed of threads), in each of the experiments the PaRLSched scheduler achieved similar or better average processing time per thread than the OS scheduler. This means that threads on the average did more work per unit of time under the PaRLSched scheduler. However, as can be seen from the execution times, this does not necessarily imply a shorter completion time, as there may be other factors that influence this.

5.2 Thread pinning and memory binding

In this last experiment, the PaRLSched scheduler was extended with the memory binding of a newly allocated (*stack*) memory to the selected NUMA node of the running thread. The intention here is to constrain any newly allocated memory into the selected NUMA node, which may potentially further increase the running speed of the overall application. We further introduce the variable $\zeta \in [0, 1]$ that captures the minimum percentage of occupancy (among threads) requested before binding the memory of a thread into a NUMA node. For example, if $\zeta = 1/2$, it implies that a thread’s memory will be bound to a NUMA node if and only if more than 1/2 of the threads also run on that node. Intuitively, the more threads occupy a NUMA node, the more likely it is that the larger part of the shared memory will be (or should be) attached on that node. Thus, essentially, through the introduction of ζ , we get an additional control variable that may potentially affect the speed of the overall application. This is indeed the case as depicted in the statistical analysis of Table 3. It is observed that under $\zeta = 1/2$, a small decrease is observed in the completion time of the overall application (which is not observed under $\zeta = 0$).

OS		PaRLSched			PaRLSched ($\zeta = 1/2$)			PaRLSched ($\zeta = 0$)		
Mean	Dev	Mean	Dev	Diff (%)	Mean	Dev	Diff (%)	Mean	Dev	Diff (%)
1317.59	11.11	1227.89	5.61	+6.80	1223.40	4.08	+7.14	1226.88	3.36	+6.88

Table 3. Statistical results regarding the completion time (in *sec*) of OS and PaRLSched under Experiment Group D ($\epsilon = 0.3/\bar{v}_i/10^8$, $\lambda = 0.1/\bar{v}_i/10^8$).

6 Conclusions and future work

We proposed a measurement- (or performance-) based learning scheme for addressing the problem of efficient dynamic pinning of parallelized applications into many-core systems under a NUMA architecture. According to this scheme, a centralized objective is decomposed into thread-based objectives, where each thread is assigned its own utility function. Allocation decisions were organized into a hierarchical decision structure: at the first level, decisions are taken with respect to the assigned NUMA node, while at the second level, decisions are taken with respect to the assigned CPU core (within the selected NUMA node). The proposed framework is flexible enough to accommodate a large set of actuation decisions, including memory placement. Moreover, we introduced a novel learning-based optimization scheme that is more appropriate for administering actuation decisions under a NUMA architecture, since a) it provides better control over the switching frequency and b) it provides better adaptivity to variations in the performance, since the experimentation probability is directly influenced by the current performance.

We demonstrated the utility of the proposed framework in the maximization of the running average processing speed of the threads. Through experiments, we observed that the PaRLSched dynamic scheduler can ensure that the running average speed of the parallelized application will be either larger than or equal to the corresponding speed under the OS's scheduler. In some cases, this also corresponded to the decrease in the execution time of an application by up to 16%.

For future work, we plan to extend this work by extending our heuristics with mechanisms to assign a group of processors to a single thread, instead of just one. This will allow us to take advantage of the internal load balancing that the Linux operating system does, thus eliminating the cases where the OS outperforms our scheduler due to its rigidity in terms of the placement of threads. Additionally, we will evaluate our scheduler on a wider range of use cases and under more variable conditions, in terms of the type and amount of external load, to determine the exact settings in which it can bring benefits to the execution time, compared to the Linux scheduler.

References

- [1] Angelis, F.D., Boaro, M., Fuselli, D., Squartini, S., Piazza, F., Wei, Q.: Optimal home energy management under dynamic electrical and thermal constraints. *IEEE Transactions on Industrial Informatics* 9(3), 1518–1527 (Aug 2013)
- [2] Bini, E., Buttazzo, G.C., Eker, J., Schorr, S., Guerra, R., Fohler, G., Årzén, K.E., Vanessa, R., Scordino, C.: Resource management on multicore systems: The ACTORS approach. *IEEE Micro* 31(3), 72–81 (2011)
- [3] Brecht, T.: On the importance of parallel application placement in NUMA multiprocessors. In: *Proceedings of the Symposium on Experiences with*

- Distributed and Multiprocessor Systems (SEDMS IV). pp. 1–18. San Deigo, CA (Jul 1993)
- [4] Broquedis, F., Furmento, N., Goglin, B., Wacrenier, P.A., Namyst, R.: ForestGOMP: An efficient OpenMP environment for NUMA architectures. *International Journal Parallel Programming* 38, 418–439 (2010)
 - [5] Chasparis, G.C., Maggio, M., Bini, E., Ārżén, K.E.: Design and implementation of distributed resource management for time-sensitive applications. *Automatica* 64, 44–53 (2016)
 - [6] Chasparis, G.C., Rossbory, M.: Efficient Dynamic Pinning of Parallelized Applications by Distributed Reinforcement Learning. *International Journal of Parallel Programming* pp. 1–15 (Nov 2017), <https://link.springer.com/article/10.1007/s10766-017-0541-y>
 - [7] Dorigo, M., Stützle, T.: *Ant Colony Optimization*. Bradford Company, Scituate, MA, USA (2004)
 - [8] Inaltekin, H., Wicker, S.: A one-shot random access game for wireless networks. In: *International Conference on Wireless Networks, Communications and Mobile Computing* (2005)
 - [9] Klug, T., Ott, M., Weidendorfer, J., Trinitis, C.: *autopin* - automated optimization of thread-to-core pinning on multicore systems. In: Stenstrom, P. (ed.) *Transactions on High-Performance Embedded Architectures and Compilers III, Lecture Notes in Computer Science*, vol. 6590, pp. 219–235. Springer Berlin Heidelberg (2011)
 - [10] Mucci, P.J., Browne, S., Deane, C., Ho, G.: PAPI: A portable interface to hardware performance counters. In: *Proceedings of the Department of Defense HPCMP Users Group Conference*. pp. 7–10 (1999)
 - [11] Olivier, S., Porterfield, A., Wheeler, K.: Scheduling task parallelism on multi-socket multicore systems. In: *ROSS'11*. pp. 49–56. Tuscon, Arizona, USA (2011)
 - [12] Rivera, F.F., Pena, T.F., Cabaleiro, J.C. (eds.): *Efficient Dynamic Pinning of Parallelized Applications by Reinforcement Learning with Applications*, *Lecture Notes in Computer Science*, vol. 10417. Springer International Publishing (2017), <http://link.springer.com/10.1007/978-3-319-64203-1>, dOI: 10.1007/978-3-319-64203-1
 - [13] Subrata, R., Zomaya, A.Y., Landfeldt, B.: A cooperative game framework for QoS guided job allocation schemes in grids. *IEEE Transactions on Computers* 57(10), 1413–1422 (Oct 2008)
 - [14] Tembine, H., Altman, E., ElAzouri, R., Hayel, Y.: Correlated evolutionary stable strategies in random medium access control. In: *International Conference on Game Theory for Networks*. pp. 212–221 (2009)
 - [15] Wei, G., Vasilakos, A.V., Zheng, Y., Xiong, N.: A game-theoretic method of fair resource allocation for cloud computing services. *The Journal of Supercomputing* 54(2), 252–269 (Nov 2010)